Correlation-Based Method for Sentiment Classification

Hussam Hamdan
Laboratoire d'Informatique de Paris 6 (LIP6), Pierre and Marie Curie University
UMR 7606, 4 place Jussieu, 75005, Paris, France
Hussam.Hamdan@lip6.fr

arXiv:1610.03120v1 [cs.CL] 10 Oct 2016

Abstract

The classic supervised classification algorithms are efficient, but time-consuming, complicated and not interpretable, which makes it difficult to analyze their results that limits the possibility to improve them based on real observations. In this paper, we propose a new and a simple classifier to predict a sentiment label of a short text. This model keeps the capacity of human interpret-ability and can be extended to integrate NLP techniques in a more interpretable way. Our model is based on a correlation metric which measures the degree of association between a sentiment label and a word. Ten correlation metrics are proposed and evaluated intrinsically. And then a classifier based on each metric is proposed, evaluated and compared to the classic classification algorithms which have proved their performance in many studies. Our model outperforms these algorithms with several correlation metrics.

1 Introduction

Polarity classification is the main task of Sentiment Analysis. The interest in polarity detection has more and more increased with the growing of user- generated data which may be interesting for many who want to understand the opinions of the users or customers towards an issue or a service.

Much research has been done on the document level sentiment analysis, but recently the focus has been on the short text which can be a sentence or a clause. For example, the sentiment analysis in Twitter where the tweet cannot exceed 140 characters, and also in on-line reviews where we seek to extract the polarity towards an aspect or a topic of interest. In SemEval workshops, many tasks have been proposed to deal with the polarity detection in short texts. Sentiment Analysis in Twitter has started since 2013 where the participants have been asked to give a polarity label to a tweet or to a topic discussed within the tweet. Aspect Based Sentiment Analysis task has also started since 2014 where one of its objectives is to give a sentiment label to a known aspect of restaurant or laptop.

On-line Short text is different from long text. The main characteristic is that the length is very short, no longer than 140 characters in Twitter; it also suffers from Sparseness, it contains several words which don not provide enough words co-occurrence. In addition to many misspellings, non-standard terms and noise.

Many supervised classification algorithms have been used to classify a short text, such as k-nearest neighbors (k-NN), Naive-Bayes, Maximum Entropy (ME), Logistic Regression (LR) and Support Vector Machines (SVM). The simple algorithms like Naive-Bayes and the more complicated ones like SVM and LR seem to give fair results in many studies. These algorithms try to assign a weight for each feature, then to predict the new text according to its final score, the process of weighting computing is time-consuming and their resulting weights are not understandable which makes it difficult for human interpret-ability and limits the ability to extend and integrate other techniques into the model. Normally, a hand-crafted feature extraction is done to improve the performance, but in an intuitive way, the weights assigned to some features can not be interpretable in many cases, therefore if we want to extract some features to improve the performance we cannot guarantee the effect of the new features.

In other hand, the results of some unsupervised sentiment classifiers can be human interpretable. For example, some unsupervised methods compute the Pair-Wise Mutual Information PMI between the phrase and some positive and negative

indicators to attribute the text to the most associated sentiment label.

In this study, we propose to use and extend the methods which have been used for unsupervised sentiment classification but in a supervised manner, thus instead of calculating the word association with the sentiment labels using some seed positive and negative words on a big un labelled corpus, we propose to use the labeled corpus for computing the word associations in order to predict the sentiment label, in this way if the classifier gives a wrong results we can understand the reason of the wrong decision and improve the performance in a wise way. Therefor, our objective is to build a simple supervised classifier which gives a good performance, but keeps the results understanding by human to integrate other techniques, rules and real observations. In this paper, we study several correlation metrics which compute the association between a term and each sentiment label using labeled data, then we use these scores to predict the polarity of a new text. The final classification model is simple, it just sums the scores given by a correlation-base metric for each sentiment label and then assign the text to the sentiment label that gets the highest score.

The paper is organized as follows. We begin with a description of related work in Section 2. Next, we describe the correlation metrics in Section 3. The proposed model is presented in Section 4. The data sets, experiments and evaluations will be discussed in Section 5. The conclusion and future work in Section 6.

2 Related Work

The correlation metrics have been widely used in supervised text classification for feature selection. In sentiment Analysis, some metrics have been used in two scenarios: (1) to measure the prior polarity of the words in order to use it for sentiment classification or sentiment lexicon construction, (2) to compute a term weight in order to integrate it in document representation for SVM classifier. The first scenario has been used in unsupervised context, but the second is supervised.

The early work (Turney, 2002) estimated the sentiment orientation (SO) of the extracted phrases using the Pointwise Mutual Information (PMI). The sentiment orientation of a phrase is computed based on its association with the positive reference word ”excellent” and the negative reference word ”poor”. Authors in (Turney and Littman, 2003) used SO to compute the sentiment orientation of a given word. They computed the orientation of the word from the strength of its association with a set of positive words (good, nice, excellent, positive,fortunate, correct, and superior), minus the strength of its association with a set of negative words (bad, nasty, poor, negative, unfortunate, wrong, and inferior). In (Mohammad, 2012) Authors collected a set of 775,000 tweets to generate a large word-sentiment association lexicon; a tweet was considered positive if it has one of 32 positive hashtagged seed words, and negative if it has one of 36 negative hashtagged seed words; the association score for a term was calculated using SO. Authors in (Mohammad et al., 2013) used similar method on the sentiment140 corpus (Go et al., 2009), a collection of 1.6 million tweets that contains positive and negative emoticons; the tweets are labeled positive or negative according to the emoticons. I In the second scenario, the researchers have used the correlation metrics to weighting the terms. While (Pang et al., 2002) reported that binary weight schema provides good accuracy with SVM, recent research has focused on more complex weighting schema which called supervised weighting metrics ,as they exploit the categorical information, or correlation metrics as they measure the correlation between a word and a category or class. Some metrics have been adopted from information retrieval such as DelatIDF (Martineau and Finin, 2009) (Paltoglou and Thelwall, 2010), later on several metrics have been proposed involving those adopted from information theory and widely used in text classification such as IG (information gain), MI (Mutual Information) (Deng et al., 2014). Recently, Wu and Gu (Wu and Gu, 2014) also tested several methods adopted from information retreival and information theory, they showed that existing methods suffer from the problem of over-weighting and introduced normalized formula of some existing metrics. They also proposed a new metric called natural entropy (ne) also inspired from information theory. Some research has used some metrics as Z score for feature extraction as in (Hamdan et al., 2014).

3 Correlation Metrics

The correlation metrics measure the degree of the association of the term with positive, negative and

neutral sentiment. These metrics are supervised, they need a training documents to compute the association between terms and sentiment labels. Before describing the ten supervised metrics used in this study, we present the notation used to define these metrics, we use the words term and feature interchangeably:
N : Number of documents in the corpus;
N_c : Number of documents in the class c;
$N_{\bar{c}}$: Number of documents out of the class c;
df : Number of documents containing the feature f in whole corpus;
df_c : Number of documents in class c containing the feature f;
$df_{\bar{c}}$: Number of documents out of class c containing the feature f;
$p(c)$:The probability of the class c;
$p(c|f)$: The probability of class c given the feature f;
$p(f)$:The probability of the feature f in the corpus;
$d\bar{f}$: The number of documents which don't contain the feature f in the corpus;
$p(\bar{c}|f)$:The complimentary probability of class c given the feature f;
$p(c, f)$The joint probability of the class c and the feature f;
$p(c, \bar{f})$: The probability of class c in the documents which don't contain the feature f;
$d\bar{f}\bar{c}$:The number of document out of class c and don't contain the feature f;
$d\bar{f}c$: The number of document don't contain the feature f in the class c.
The Ten weighting metrics are:

Natural Entropy (ne)
The more uneven the distribution of documents where a feature occurs, the larger the weight of this feature is. Thus, the entropy of the feature can express the uncertainty of the class given the feature. One minus this degree of uncertainty boosts the features that unevenly distributed between the category and other categories.

$$ne(f,c) = 1 + (p(c|f).log(p(c|f)) + p(\bar{c}|f).log(p(\bar{c}|f)))$$

This metrics was firstly proposed in (Wu and Gu, 2014), ne score is always between 0 and 1, and it assigns a high score for the words unevenly distributed over the classes, but it cannot discriminate the positive words from the negative ones. Therefore, authors in (Hamdan et al., 2015) have used the df_c and $df_{\bar{c}}$ for discriminating the positive words from the negative ones; if $df_c > df_{\bar{c}}$ then the word is considered positive else it is considered negative. In this work a negative value is given for a negative word.

Pairwise Mutual Information (pmi)
PMI is a measure of association used in information theory and statistics (Church and Hanks, 1990 03).

$$pmi(f,c) = log(\frac{p(c,f)}{p(c).p(f)})$$

Odds Ratio (orr)
Orr gives a positive score to features that occur more often in one category than in the other, and a negative score if it occurs more in the other. A score of zero means the the odds for a feature to occur in one category is exactly the same as the odds for it to occur in the other (Shaw, 1995).

$$orr(f,c) = log(\frac{p(f|c).(1-p(f|\bar{c}))}{p(f|\bar{c}).(1-p(f|c))})$$

Categorical Proportional Difference (cpd)
Cpd is a ratio that considers the number of documents of a category in which the feature occurs and the number of documents from other categories in which the feature also occurs (Simeon and Hilderman, 2008).

$$cpd(f,c) = \frac{df_c - df_{\bar{c}}}{df}$$

Kullback-Leibler Divergence (kl)
Kl is a non-symmetric measure of the difference between the distribution of the category and the distribution of the category given the feature. A measure of how dissimilar the two distributions are. Useful feature value imply a high degree of dissimilarity.

$$kl(f,c) = p(c|f).log(\frac{p(c|f)}{p(c)})$$

Relevance Frequency (rf)
The basic idea of rf is to boost the higher frequency terms in the positive category than in the negative one, that helps in selecting the positive samples from the negative ones (Lan et al., 2009).

$$rf(f,c) = log(2 + \frac{dfc}{max(1, df_{\bar{c}})})$$

Multinomial Z Score (zd)
Suppose that the feature follows multinomial distribution over the classes, zd calculates Z transformation for a feature in each class, zd boosts the highly unevenly distributed features among the classes, it gives high positive score for a feature in the class where it is highly frequent and negative score in the class where it rarely appears (Hamdan et al., 2014).

$$zd(f,c) = \frac{dfc - p(f).N_c}{\sqrt{N_c.p(f).(1-p(f))}}$$

Delta BM25 IDF (dbidf)
Dbidf is a variant of delta-idf measure, BM25 idf variant is used instead of classical idf. dbidf was published in Sentiment Analysis literature (Paltoglou and Thelwall, 2010).

$$dbidf(f,c) = log(\frac{(N_c - df_c + 0.5).df_{\bar{c}} + 0.5}{(N_{\bar{c}} - df_{\bar{c}} + 0.5).df_c + 0.5})$$

Weighted Log Likelihood Ratio(wllr)
Wllr is a measure of how dissimilar the distribution of the feature given the category and the distribution of the feature given the other categories (Nigam et al., 2000).

$$wllr(f,c) = p(f|c).log(\frac{p(f|c)}{p(f|\bar{c})})$$

NGL Coefficient (ngl)
Ngl is a variant of the Chi square metric.It measures the lack of independence between the feature and the category. The higher value, the closer relationship the feature and the class have (Ng et al., 1997).

$$ngl(f,c) = \frac{\sqrt{N}(df_c.d\bar{f}_{\bar{c}} - df_{\bar{c}}.d\bar{f}_c)}{\sqrt{df.d\bar{f}.Nc.N\bar{c}}}$$

4 Correlation-Based Model

In this section, we introduce our model which directly relates each short text instance to a sentiment label. How a short text instance, denoted by s, is related to a sentiment label sl can be measured based on the correlation of the overlapping words among the words in the training set and the sentence s. Thus, for each correlation metric which measures the degree of association of each word to each sentiment label, we compute the final polarity score of a short text as the sum score of its words. Three sum scores are computed, one for each sentiment label, then the sentence will be assigned to the sentiment label which has the highest sum score. For calculating the sum score of a sentence s in a sentiment label or class i, we sum the scores of the sentence words.
The correlation between s and sl is calculated based on scores as:

$$sumscore(s,i) = \sum_{t \in s} metricScore(t,i)$$

metricScore(t,i) stands for the score of correlation between the term t and the sentiment label i, this score is given by any of the ten correlation metrics presented in the previous section. Finally, s is assigned to a sentiment label which have the highest score. For example, let's assume that s=" happy day" the PMi metric gives three scores for the word happy 0,0.2,1.2 in negative, neutral and positive class respectively and the three scores for day 0.3,0.9,0.2. The model produce three score, one for each sentiment label, positive score=1.2+0.2=1.4, negative score=0+0.3=0.3, neutral score=0.2+0.9=1.1 while the positive score is the highest one. Therefore, the sentence is classified as positive.

5 Experiments and Evaluations

For testing our proposed model, we conduct two types of evaluations: intrinsic and extrinsic. For intrinsic evaluation, we test the performance of the proposed correlation metrics, therefore we compare the degree of prior polarity with a manually annotated list of Twitter terms. For testing our model, an extrinsic evaluation is done using two short text data set, where we calculate the correlation scores based on the training set of each data set.

5.1 Training and Testing Data

We have used two data sets, the first one is extracted from Twitter which has been provided in SemEval 2013 for subtask B of sentiment analysis in Twitter (Nakov et al., 2013). The participants have been provided with training tweets annotated positive, negative or neutral. We downloaded these tweets using the given script. We obtained 9646 tweets, the whole training data set is used for training. The test data set provided in SemEval-2015 containing about 2390 tweets (Rosenthal et al., 2015) is used for evaluating our system.
The second data set is extracted from Laptop re-

Data	All	Positive	Negative	Neutral
Twitter				
train	9684	3640	1458	4586
test	2390	1038	365	987
Laptop				
train	3989	277	1739	1973
test	1883	173	761	949

Table 1: Sentiment labels distribution in the training and testing data sets in Twitter and Laptop reviews.

views, provided by SemEval 2015 ABSA organizers (Pontiki et al., 2015) where each review is composed of several sentences, the participants are asked to detect the polarity of each senetce. Table 1 shows the distribution of each label in each data set.

5.2 Correlation Metrics Evaluation

This evaluation is done for measuring the performance of the correlation metrics in determining strength of association of Twitter terms with positive sentiment. A ranked list of twitter terms obtained from human annotations was provided by SemEval-2015 organizers (Rosenthal et al., 2015), the original ranked list contains 1315 terms, we used only 552 terms which exist in our training data set. We first compute the score of each term as the score of correlation with positive sentiment minus the score of correlation with the negative and neural ones, and then comparing this ranked list to the human annotated one. Kendall's Tau and Spearman will be used as metrics to compare the ranked lists. Table 2 shows the results for each metric with Kendall metric and Spearman. Kl gives the highest correlation 64% followed by cpd and rf 44% then orr metric with 43% with Kendall's Tau.

Note that several metrics give the same results such ne, pmi, zd. Therefore, we measure the correlation between each pairs of metrics to discover how the metrics which based on different mathematical calculations can give different scores for each word but maintains the same ranking. Table 3 shows the correlation between each two pairs. Note that the metrics which give the same kendall score are so correlated as zd,ngl and ne,pmi,orr,kl. In other way, it seems that the metrics constitute some groups, each group tend to give the same ranking.

metric	Kendall	Spearman
ne	0.42	0.58
pmi	0.42	0.58
orr	0.43	0.6
cpd	0.44	0.61
kl	0.46	0.63
rf	0.44	0.6
dbidf	0.43	0.6
zd	0.42	0.59
wllr	0.38	0.54
ngl	0.41	0.57

Table 2: The correlation between the scores given by each correlation metric and the manually annotated ranked Twitter terms list.

5.3 Correlation-Based Model Evaluation

For evaluation our correlation-based model Section 3, we used two data sets extracted from Twitter and Laptop reviews. We have tested the model using one of each correlation metrics in addition to the evaluation using the SVM classifier and Logistic Regression which have achieved a robust baseline in SemEal2015 competition.

Table 4 shows the results using our proposed model with the ten metrics, SVM and LR, the second column contains the f-score of positive and negatives sentiment labels, the metric of SemEval evaluation. But the third column contains the accuracy in Laptop reviews data set which was the standard evaluation metric in SemEval-2015 for this task.

model	Twitter	Laptop
SVM	52.35	69.97 (664/949)
LR	53.38	71.65 (680/949)
ne	42.26	72.29 (686/949)
pmi	53.15	71.13 (675/949
orr	54.11	72.4 (687/949)
cpd	19.20	67.65 (642/949)
kl	57.30	69.02 (655/949)
rf	24.46	71.65 (680/949
dbidf	13.25	8.32 (79/949)
zd	49.24	69.65 (661/949
wllr	45.11	50.68 (481/949)
ngl	50.27	69.65 (661/949

Table 4: The results given by SVM, LR and our proposed model with one correlation metric each time. The f1-score of negative and positive labels is used in Twitter set but the accuracy metric in Laptop set.

metric	ne	pmi	orr	cpd	kl	rf	dbidf	zd	wllr	ngl
ne	1.0	0.99	0.98	0.91	0.85	0.85	0.98	0.87	0.77	0.86
pmi	0.99	1.0	0.99	0.91	0.84	0.85	0.99	0.87	0.76	0.86
orr	0.98	0.99	1.0	0.96	0.9	0.91	1.0	0.91	0.8	0.91
cpd	0.91	0.91	0.96	1.0	0.96	0.99	0.95	0.93	0.83	0.94
kl	0.85	0.84	0.9	0.96	1.0	0.96	0.89	0.92	0.87	0.92
rf	0.85	0.85	0.91	0.99	0.96	1.0	0.91	0.91	0.82	0.93
dbidf	0.98	0.99	1.0	0.95	0.89	0.91	1.0	0.9	0.8	0.9
zd	0.87	0.87	0.91	0.93	0.92	0.91	0.9	1.0	0.96	1.0
wllr	0.77	0.76	0.8	0.83	0.87	0.82	0.8	0.96	1.0	0.95
ngl	0.86	0.86	0.91	0.94	0.92	0.93	0.9	1.0	0.95	1.0

Table 3: The correlation between each two pairs of the ten correlation metrics using Kendall'Tau measure.

In Twitter set, the performance of kl is the best, it achieves 57.30% with a gain of 4.95% over SVM and 3.92% over LR, note that kl was the best metric in the metric evaluation, kl gives a significant improvement of SVm and LR. orr metric gives a better performance over SVM and LR while pmi outperforms SVM but not LR. Other metrics are under the baseline of SVM and LR.

Figure 1 demonstrates the distribution of words in Twitter set in positive class using the ten correlation metrics, each metric distributes the words in a different way. While some metrics as dbidf gives good ranking score according to the metric evaluation, it is bad in model evaluation, the reason may be that the scores given by dbidf have a wide range which may negatively affect the sum of scores. We also remark that some metrics produce some small negative scores such kl but others produce more large negatives ones such zd, pmi, ngl. In our experiments we choose a threshold to be 0 for kl because its negative values so small to be considered but the other metrics is evaluated without such threshold. It should note that the seeking for an appropriate threshold for each metric may lead to more robust model.

The third column in Table 4 is the accuracy score as it has been considered as evaluation measure in SemeEval-2015, the accuracy is followed by the number of correct sentences over the total number of sentences. The results show that orr metric outperforms other metrics, it achieves 72.4% with 2.43% over SVM and 0.75% over LR. ne metric also outperforms SVM and LR, rf is similar to LR and pmi gives a fair result but lower than SVM nad LR. kl metric which is the best one for Twitter set seems not efficient in Laptop and in general our model is not so as efficient in Laptop set as in twitter data set. One possible reason that the metrics need a sufficient amount of training data to produce a reliable scores.

Therefore, the correlation based methods seem to be promising in the sentiment analysis of short text. They are simple, fast and interpretable methods. Some metrics like orr and pmi give good results in both data sets and outperforms the two classic models SVM and LR. In average the orr metric seems to be the best metric.

Table 5 shows the ten high words produced for the best metrics with there scores. Some metrics produce the same ranked list but with different scores. For example, kl, pmi and orr produce the same list but with different scores, zd and ngl do the same. Therefore, the proposed model can be modified to take into account this observation, a normalization factor can be added to make the metrics which produce the same rank give the same results. A deeper research may led to classify the metrics into some groups where each group produces the same ranked list and a combination of different groups may improve the classification results.

6 Conclusion

In this paper, we presented ten correlation metrics for measuring the correlation between the words and the sentiment labels. Then, we proposed a classification model based on a correlation metric. We have evaluated the metrics intrinsically using a ranked list of twitter terms, and the proposed model was evaluated in sentiment classification task using two short text data sets extracted from twitter and laptop reviews. The proposed

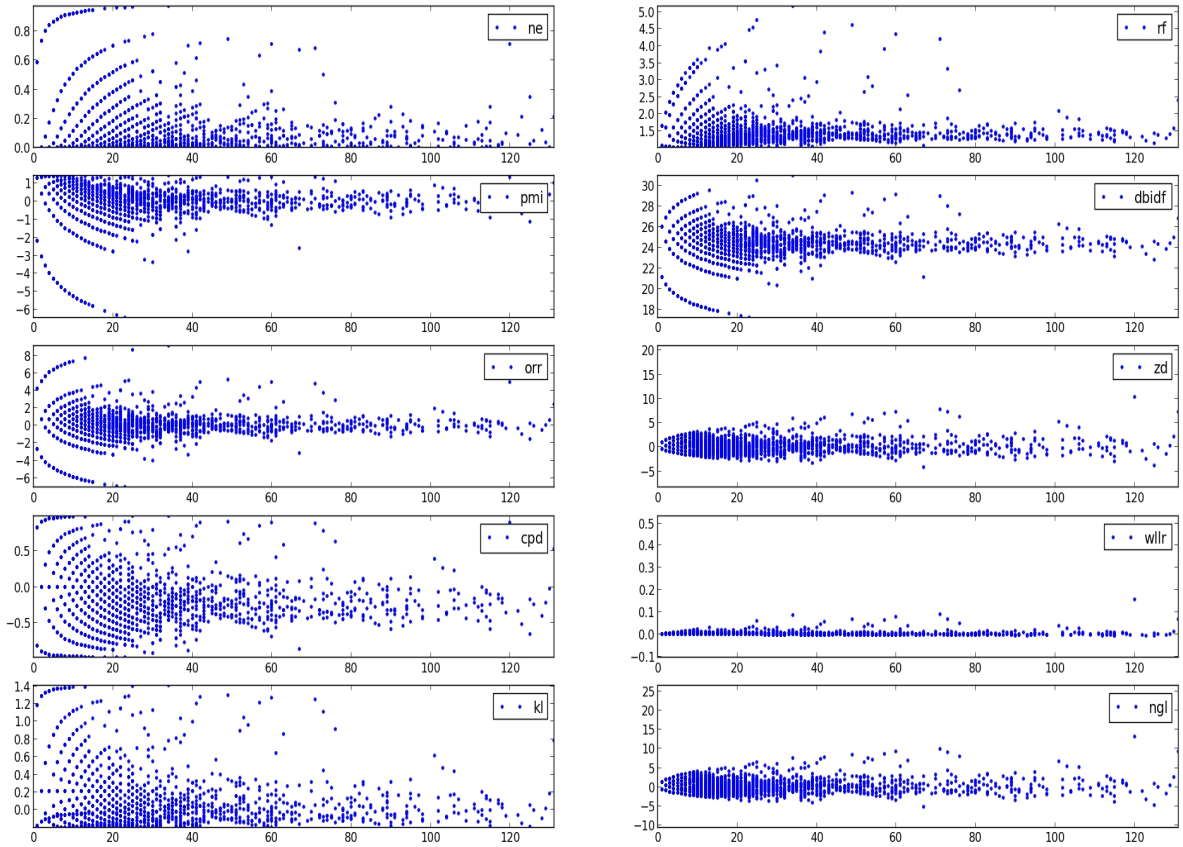

Figure 1: The distribution of each correlation metric in Twitter data set. X-axis represents the frequency, y-axis represents the score of each word given by the metric. Words which have the same frequency in the corpus, have different scores according to each metric because they have different distributions over the sentiment labels.

ngl		zd		kl		orr		pmi	
!	26.44	!	20.89	glad	1.4	glad	9.16	glad	1.41
good	26.28	good	14.8	congrats	1.4	congrats	8.71	congrats	1.41
:)	18.73	:)	14.3	delicious	1.39	delicious	7.77	delicious	1.4
love	18.11	love	13.21	hilarious	1.38	hilarious	7.39	hilarious	1.4
happy	16.72	happy	12.66	stoked	1.38	stoked	7.24	stoked	1.4
great	16.02	great	11.64	bacolod	1.38	bacolod	7.24	bacolod	1.4
excited	14.73	excited	10.32	jennette	1.38	jennette	7.24	jennette	1.4
wait	13.06	wait	9.87	celebrated	1.38	celebrated	7.07	celebrated	1.39
best	12.5	best	9.79	enjoyed	1.38	enjoyed	7.07	enjoyed	1.39
fun	12.39	fun	7.83	wonderful	1.38	wonderful	7.07	wonderful	1.39

Table 5: The top ten words given by ngl, zd, kl and pmi correlation metrics with their scores in Twitter set.

model seem to give good results and outperforms some classic supervised algorithms with some correlation metrics. The results given by the model can be interpretable and analyzed in order to improve it in a wise way. In the next work, we plan to explore how we can combine the NLP techniques in our model to take into account the negation, the reversal words, and the modifiers. We also plan to combine noisy labeled or unlabeled data which may help to get more robust correlation scores.

References

[Church and Hanks1990 03] Kenneth Ward Church and Patrick Hanks. 1990-03. Word association norms, mutual information, and lexicography. 16(1):22–29.

[Deng et al.2014] Zhi-Hong Deng, Kun-Hu Luo, and Hong-Liang Yu. 2014. A study of supervised term weighting scheme for sentiment analysis. 41(7):3506 – 3513.

[Go et al.2009] Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. pages 1–6.

[Hamdan et al.2014] Hussam Hamdan, Patrice Bellot, and Frederic Bechet. 2014. The impact of z_score on twitter sentiment analysis. In *In Proceedings of the Eighth International Workshop on Semantic Evaluation (SemEval 2014)*, page 636.

[Hamdan et al.2015] Hussam Hamdan, Patrice Bellot, and Frederic Bechet. 2015. Sentiment lexicon-based features for sentiment analysis in short text. In *In Proceeding of the 16th International Conference on Intelligent Text Processing and Computational Linguistics*.

[Lan et al.2009] Man Lan, Chew Lim Tan, Jian Su, and Yue Lu. 2009. Supervised and traditional term weighting methods for automatic text categorization. 31(4):721–735.

[Martineau and Finin2009] Justin Martineau and Tim Finin. 2009. Delta TFIDF: An improved feature space for sentiment analysis. In *ICWSM*.

[Mohammad et al.2013] Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRCCanada: Building the state-of-the-art in sentiment analysis of tweets. In *In Proceedings of the International Workshop on Semantic Evaluation, SemEval 13*.

[Mohammad2012] Saif Mohammad. 2012. #emotional tweets. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 246–255. Association for Computational Linguistics.

[Nakov et al.2013] Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013. SemEval-2013 task 2: Sentiment analysis in twitter. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320. Association for Computational Linguistics.

[Ng et al.1997] Hwee Tou Ng, Wei Boon Goh, and Kok Leong Low. 1997. Feature selection, perceptron learning, and a usability case study for text categorization. In *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '97, pages 67–73. ACM.

[Nigam et al.2000] Kamal Nigam, Andrew Kachites McCallum, Sebastian Thrun, and Tom Mitchell. 2000. Text classification from labeled and unlabeled documents using EM. 39(2):103–134.

[Paltoglou and Thelwall2010] Georgios Paltoglou and Mike Thelwall. 2010. A study of information retrieval weighting schemes for sentiment analysis. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 1386–1395. Association for Computational Linguistics.

[Pang et al.2002] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10*, EMNLP '02, pages 79–86. Association for Computational Linguistics.

[Pontiki et al.2015] Maria Pontiki, Dimitrios Galanis, Haris Papageogiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. SemEval-2015 task 12: Aspect based sentiment analysis. In *In Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*.

[Rosenthal et al.2015] Sara Rosenthal, Preslav Nakov, Svetlana Kiritchenko, Saif M. Mohammad, Alan Ritter, and Veselin Stoyanov. 2015. SemEval-2015 task 10: Sentiment analysis in twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation*, SemEval '2015. Association for Computational Linguistics.

[Shaw1995] W. M. Shaw, Jr. 1995. Term-relevance computations and perfect retrieval performance. 31(4):491–498.

[Simeon and Hilderman2008] Mondelle Simeon and Robert Hilderman. 2008. Categorical proportional difference: A feature selection method for text categorization. In *Proceedings of the 7th Australasian Data Mining Conference - Volume 87*, AusDM '08, pages 201–208. Australian Computer Society, Inc.

[Turney and Littman2003] Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. 21(4):315–346.

[Turney2002] Peter D. Turney. 2002. Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 417–424. Association for Computational Linguistics.

[Wu and Gu2014] Haibing Wu and Xiaodong Gu. 2014. Reducing over-weighting in supervised term weighting for sentiment analysis. In *COLING 2014, 25th International Conference on Computational*

Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland, pages 1322–1330.